Titus Winters - Research Statement

My research interests fall in two main subject areas: system security and education. At first glance these seem to be two very disparate areas. In fact, my interest in both subjects is strongly supported by my desire to develop high-impact applications. In security I build tools for information gathering to better know the threats that we face on the Internet today. In education I build tools to reduce instructor workloads or give students deeper understanding of the subject matter. I also develop statistical methods to provide instructors greater insight into student learning. My interests in security and education were both fostered while I was an undergraduate at Harvey Mudd College, and have been reinforced during my graduate work at the University of California at Riverside.

My focus in networking and security stems directly from my work in the Aerospace Corporation's Trusted Systems Lab, where I worked from 2000-2002. There I focused on honeypots and decoying technologies, under funding from the National Security Agency. Honeypots are systems placed on a network with the intent that they will be attacked, in the hopes of monitoring attackers in the act. My technical responsibilities ranged from demonstrating flaws in commercial honeypots to developing prototype security systems, as well as recruiting for summer positions and developing proposals for NSA-funded projects.

I worked on two main projects at Aerospace: Melli, a Linux kernel module honeypot, and DRAGON, a prototype for a new intrusion-reaction paradigm. Melli was designed to function as a normal Linux installation, but secretly log all activity on the system, and protect critical system resources (log files, etc) from attackers. The concept was for Melli to be deployed on a public network and wait for an attacker to access the system. Once compromised, Melli would log all of the attacker's activities without their knowledge, and would itself be effectively undetectable. DRAGON employed similar decoying techniques: DRAGON was a prototype for an Intrusion Deflection System, or IDFS. DRAGON was composed of a router that we developed, a decoy Melli host, and communications channels for receiving information from an Intrusion Detection System (IDS). When an IDS alert was sent to DRAGON, the router would identify all of the communications coming from the source of the attack, and would send necessary information about those connections to the decoy host. All current and future communications would then be transparently redirected away from the original target and into the decoy host. This would allow for the monitoring of much more advanced attackers, i.e. those attacking specific hosts rather than random addresses. Both of these projects were unclassified work funded by the NSA.

My senior project at Harvey Mudd grew directly out of my work on DRAGON. While developing DRAGON it became clear to me that much of the software I was developing for packet handling on the router could be easily adapted for building a hands-on laboratory for teaching low-level networking, letting students build their own network stack. I proposed building such a system to Professor Mike Erlinger, and TinkerNet was started in 2001. The project continues to this day under a national deployment grant from NSF. TinkerNet was an excellent experience for me in terms of an introduction to building good user interfaces, deploying a complex software system, and especially for building educational tools.

This interest in building educational tools carried over into my work at UCR. Here I have focused on several large projects in two main areas: instructional tools and machine-learning applied to educational data gathered by those tools. In the spring of my first year at UCR, I began working with Professor Tom Payne. We brainstormed for weeks about software systems that could track student abilities in every topic in the curriculum, and what types of tools would be needed for such a project. The tools that were needed were Agar and HOMR, two projects I developed and managed along with a volunteer group of undergraduate students. Agar is a Computer Aided

Assessment (CAA) framework, allowing for automated grading of programming assignments, with easy extensibility for new methods of functional testing. Unlike most CAA systems, Agar also gives significant benefit to the grading of subjective work, by allowing easy re-use of grader marks and comments to students. Since many mistakes are made repeatedly on a given assignment, writing out the feedback and assigning a penalty only once results in increased grader productivity. Future cases of the same problem can be graded by using drag-and-drop to place a copy of the comment in the later submission.

My other educational software tool is known as Homebrew Optical Mark Recognition (HOMR). HOMR allows a user to design their own OMR (fill-in-the-bubble) forms, print them on any printer, fill them in with either pencil or pen, scan them on any scanner, and get back results for each bubble in an easy-to-parse format. This system forms a cornerstone for the data collection portions of my dissertation work. HOMR relies heavily on supervised-learning techniques, classification, and computer vision. HOMR uses standard vision and classification algorithms to normalize the scans, allowing for skews, rotation, offsets of the page, extra noise in scans, staples, torn corners, etc. Once the page has been normalized, it reads in a page layout to find the normalized location of each grid of bubbles. It uses gradient ascent to center a grid around those bubbles, extracts an image sample of each bubble, and uses a boosting classifier to identify each bubble. Unlike most OMR systems, HOMR caters to the natural behavior of students by allowing bubbles to be classified as filled, empty, or crossed-out, which is graded as empty. Preliminary results indicate that HOMR can achieve error rates about 2% of those in the system used by the US Census for a similar task.

The main work of my dissertation is in processing the data provided by HOMR and Agar, focusing on the identification of questions with similar topics based on per-student scores. The major contribution of this work is a comparison of the amount of topic data present in material learned superficially versus deep knowledge, and the ability of various data mining and machine-learning algorithms to extract topics in either case. Although the results are not yet finalized, the pedagogical implications of this research are that student's ability to answer the questions in a course is not primarily based on the subject of those questions.

I am eager to continue my research in both of these subject areas. I have plans for adapting the TCP connection deflection concepts used in DRAGON for use in load-balancing, reliable communications, and host mobility applications. I hope to continue my work on HOMR to increase the accuracy and decrease the set of scanning errors that cause bad page normalization. I also look forward to adapting Agar to work on a Tablet PC platform. Continuing my dissertation work, I hope to continue my investigations into statistical methods for predicting student scores. Many of these tasks can be made concrete enough to be handed off to undergraduate researchers. As an undergraduate, I learned more in my research experiences than in my coursework, and I look forward to helping my students have similar experiences.

Viewed as a coherent whole, the thread that ties together all of my work is a need for there to be a strong potential for positive impact. Network security is a critical issue in the modern world, and becoming more relevant day by day. Tools to learn more about the black-hat hacker community are invaluable in this area. My educational interests at the theoretical/statistical level are intended to discover hidden meaning in student understanding to help instructors do a better job. The development of educational tools is aimed at freeing up instructor time. My research interests are relatively broad, but all focus on fulfilling the mission of my alma mater: to be a scientist with an awareness of society.